

REMARKS

In the final Office Action, the Examiner rejects: (i) claims 1, 5-7, 9 and 14-17 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 6,449,647 (hereinafter “Colby”) in view of newly cited U.S. Patent Publication No. 2004/0258003 (hereinafter “Kokot”); (ii) claims 2-4, 18-20 and 25 under 35 U.S.C. §103(a) as being unpatentable over Colby/Kokot in view of U.S. Patent No. 6,112,221 to (hereinafter “Bender”); (iii) claim 8 under 35 U.S.C. §103(a) as being unpatentable over Colby/Kokot in view of U.S. Patent No. 6,807,156 (hereinafter “Veres”); (iv) claims 10-12 under 35 U.S.C. §103(a) as being unpatentable over Colby/Kokot in view of U.S. Patent No. 6,981,029 to (hereinafter “Menditto”); (v) claim 13 under 35 U.S.C. §103(a) as being unpatentable over Colby/Kokot/Menditto in view of U.S. Patent No. 6,772,211 (hereinafter “Lu”); (vi) claims 21-23 under 35 U.S.C. §103(a) as being unpatentable over Colby/Kokot/Bender/Menditto; and (vii) claim 24 under 35 U.S.C. §103(a) as being unpatentable over Colby/Kokot/Bender/Menditto/Lu.

Applicants continue to respectfully traverse the §103(a) rejection on the ground that the Colby/Kokot combination fails to teach or suggest each and every limitation of the claimed invention.

To reiterate, independent claim 1 is directed to a method of processing a request to at least one server, comprising the steps of: receiving the request; and scheduling submission of the request to the at least one server based on: (i) a quality-of-service (QoS) class assigned to a client from which the request originated; (ii) a response target associated with the QoS class; and (iii) an estimated response time associated with the at least one server.

As illustrated at page 6 of the present specification, in one embodiment, the invention provides techniques for scheduling requests at back-end servers in order to provide differentiated classes of quality of service (QoS). The techniques are implemented in the form of a scheduler external to the server. Such an implementation obviates the need for any changes to the back-end server.

Colby discloses (see Abstract) a content-aware flow switch intercepts a client content request in an IP network, and transparently directs the content request to a best-fit server. The best-fit server

is chosen based on the type of content requested, the quality of service requirements implied by the content request, the degree of load on available servers, network congestion information, and the proximity of the client to available servers. The flow switch detects client-server flows based on the arrival of TCP SYNs and/or HTTP GETs from the client. The flow switch implicitly deduces the quality of service requirements of a flow based on the content of the flow.

The final Office Action cites column 2, lines 55-56, of Colby to assert that Colby discloses “scheduling submission of the request to the at least one server,” as claimed. Column 2, lines 55-56, of Colby states: “[s]pecifically, when a client in an IP network makes a content request, the request is intercepted by a content-aware flow switch, which seamlessly forwards the content request to a server that is well-suited to serve the content request.” That is, the Colby flow switch receives the request and determines the “best fit” server for handling the request. However, there is absolutely no disclosure in Colby of determining a schedule for submitting (i.e., scheduling submission) the intercepted request to this best-fit server. The claimed invention uses “(i) a quality-of-service (QoS) class assigned to a client from which the request originated; (ii) a response target associated with the QoS class; and (iii) an estimated response time associated with the at least one server” to schedule submission of the request to a server. Colby merely determines a best-fit server based on some content criteria and then forwards the request to that server.

That is, there is no notion in Colby of scheduling submissions of requests to any server whatsoever. The one and only mention of “scheduling” made in Colby appears to be at column 16, line 58, where the notion of “flow pipes” is described. As explained earlier in column 16, the content-aware flow switch of Colby can be used to front-end many web servers. Each of the physical web servers may embody one or more virtual web hosts (VWH's). Associated with each of the VWH's front-ended by the flow switch may be a “flow pipe,” which is a logical aggregation of the VWH's flows. Flow pipes guarantee an individual VWH a configurable amount of bandwidth through the content-aware flow switch. The flow switch allocates the flow pipe bandwidth and shares it among the individual flow pipes using a weighted round robin scheduling algorithm in which the weight assigned to an individual flow pipe is a percentage of the overall bandwidth available to clients. Thus, any scheduling of “flow pipes” done in Colby is restricted to this

weighted round robin approach and is not based on the three specific scheduling criteria recited in claim 1, i.e., “(i) a quality-of-service (QoS) class assigned to a client from which the request originated; (ii) a response target associated with the QoS class; and (iii) an estimated response time associated with the at least one server” to schedule submission of the request to a server.”

Unrelated to “flow pipes” and round robin schedules, the final Office Action cites column 9, lines 25-32, of Colby to reject the request submission scheduling criteria of claim 1 of “a response target associated with the QoS class.” However, all that column 9 of Colby appears to disclose is that “[i]dentifying the nature of the requested content [by the flow switch of Colby] also involves deducing, from the content request and information stored in the CSD [Content Server Database], the QoS requirements of the requested content. These QoS requirements include: Bandwidth, defined by the number of bytes of content to be transferred over the average flow duration. Delay, defined as the maximum delay suitable for retrieving particular content. Frame Loss Ratio, defined as the maximum acceptable percentage of frame loss tolerated by the particular type of content. A QoS class is assigned to a flow based on the flow’s calculated QoS requirements.” However, again, it is important to point out that any consideration of QoS requirements in Colby is to “identify the content” of the request, not to “schedule submission of the request to a server,” as is claimed.

Based on Examiner-admitted deficiencies of Colby, the Kokot reference is specifically introduced to address the failure of Colby to disclose that a quality-of-service (QoS) class is assigned to a client from which a request originated, as claimed. The final Office Action continues to cite paragraph [0116] of Kokot for support. Applicants cite below paragraph [0116] as well as the paragraphs that proceed it to illustrate why Kokot does not assign QoS classes to clients but rather packet flows (for a VoIP call) are what have QoS classes assigned thereto:

[0103] Subscriber device 118A places a VoIP call to a subscriber device 118B via network 122 by negotiating a VoIP session with subscriber device 118B. Subscriber devices 118 also send request messages to SE routers 112 requesting an enhanced QoS for the packet flow that will carry the VoIP call. In response to the request messages, SE routers 112 verify authentication and authorization information previously received from respective servers 126A and 126B to authenticate the respective subscribers, and to determine whether the respective subscribers are authorized to receive a requested

QoS class for a VoIP call. Servers 126 may be, for example, RADIUS servers. In some embodiments, SE routers 112 are served by a single server 126.

[0104] Servers 126A and 126B store QoS profiles 128A and 128B (collectively "QoS profiles 128"), respectively. QoS profiles 128 describe the QoS classes, if any, that subscribers, such as the subscribers associated with subscriber devices 118, are authorized to receive. QoS profiles 116 maintained by SE routers 112 includes QoS information received from servers 126 when subscriber associated with subscriber devices 118 previously logged on or otherwise activated their multimedia service accounts. If SE routers 112 determine that the subscribers associated with subscriber devices 118 are authorized to use the requested QoS class for the VoIP call based on information contained in QoS profiles 116, SE routers 112A and 112B identify information within QoS profiles 116A and 116B, respectively, describing packet transmission according to the QoS class requested by subscriber devices 118A and 118B for the call. QoS profiles 116 may include, for example, information describing a designated route or packet flow that may be used by SE routers 112 to provide an enhanced QoS for the VoIP call. The indicated route or packet flow may provided greater bandwidth, or may be dedicated to or provide priority for VoIP packets.

[0105] QoS profiles 116 also includes information for use by CPE devices 114 to provide an enhanced QoS for the VoIP call. SE routers 112 provide such information to CPE devices 114 to cause CPE devices to facilitate the requested QoS class for the VoIP call. CPE devices 114 store the QoS information provided by SE routers 112 as QoS profiles 120 for the layer-2 links between CPE devices 114 and subscriber devices 118. QoS profiles 120 may include, for example, information directing CPE devices 114 to provide preferential queuing for VoIP packets.

[0106] FIG. 8 is a block diagram illustrating an example SE router 112. SE router 112 can provide a requested QoS class for a packet flow, such as a VoIP call, and can also control a CPE device 114 to facilitate the requested QoS class for the packet flow, as described above. SE router 112 includes IFCs 130, inbound and outbound links 132 and 134, and a control unit 136 that maintains routing information 138 and forwarding information 140 to forward packets received on inbound links 134 as described above with reference to SE router 22, which included IFCs 50, inbound and outbound links 52 and 54, and control unit 56 that maintains routing information 58 and forwarding information 60, and FIG. 4.

[0107] Control unit 136 maintains QoS profiles 116 that are used by control unit 136 to provide a requested QoS class to one or more subscribers for one or more packet flows. Control unit 136, for example, may receive a message requesting authentication and authorization for a VoIP call with a particular QoS class from a subscriber device 118 via one of inbound links 132 and IFCs 130. Control unit 136 checks

authentication/authorization information 142 to authenticate and authorize the subscriber associated with the requesting subscriber device 118, and to access QoS information for the VoIP call stored within QoS profiles 116. As described above, the QoS profiles may include information describing a route packet flow that may be used by control unit 136 to provide packet transmission according to the requested QoS class for the VoIP call.

[0108] QoS profiles 116 may also include information indicating a queuing preference for VoIP packets to be used by a CPE device 114 to facilitate packet transmission according to the requested QoS class for the VoIP call. Control unit 136 sends a control message to the CPE device 114 via one of IFCs 130 and a respective one of outbound links 134 to direct the CPE device 114 to implement the indicated preferential queuing for the VoIP call. Control unit 136 may send control messages to the CPE device 114 on a dedicated control ATM VC, as described above with reference communication between SE routers 22 and switches 24 of FIG. 2.

[0109] Control unit 136 may include one or more microprocessors, DSPs, ASICs, FPGAs, or other logic circuitry. Control unit 136 may include memory (not shown) that stores computer-readable program instructions that cause control unit 136 to perform the functions ascribed to it herein. The memory may include any magnetic, optical, or electrical media, such as a RAM, ROM, hard disk, CD-ROM, or EEPROM. Control unit 136 may maintain routing information 138, forwarding information 140, and QoS information 116 in memory in the form of one or more tables, databases, link lists, radix trees, databases, flat files, or any other data structures.

[0110] FIG. 9 is a block diagram illustrating an example CPE device 114 that receives QoS information from a SE router 112, and dynamically configures a QoS profile 120 for a layer-2 link between CPE device 114 and a subscriber device 118 based on the QoS information. As described above, CPE device 114 may be, for example, a modem, wireless access point, or switch. CPE device 114 includes interfaces 150 for coupling CPE device 114 to one or more subscriber devices or one or more other CPE devices, and for coupling CPE device 114 to network 122, e.g., to a SE router 112 via a switch 124. Interfaces 150 may include, for example, IFCs, such as IFCs 50, 70 and 130 described above, or transceivers for communication via a wireless medium, such as communication according to one of the IEEE 802.11 family of standards. Where CPE device 114 is a modem, interfaces 150 may include or be coupled to a control unit 152 via circuitry (not shown) for modulating and demodulating signals sent or received by CPE device 114 via interfaces 150.

[0111] Control unit 152 receives cells, frames, or otherwise encapsulated packets from a switch 124, and forwards the packets therein to a connected subscriber device 118 within Ethernet frames according to either of the IEEE 802.3 or 802.11 families of

standards. Control unit 152 also receives Ethernet frames from the connected subscriber device 118, and encapsulates the packets therein for transmission to the switch 124. Control unit 152 receives a control message from SE router 112, as described above, and dynamically configures a QoS profile 120 based on the control message. Based on QoS profile 120, control unit 152 provides packet transmission on the layer-2 link according to a requested QoS class for a packet flow for the connected subscriber device 118. For example, as described above, control unit may preferentially queue packets for the packet flow based on QoS profile 120.

[0112] Control unit 152 may include one or more microprocessors, DSPs, ASICs, FPGAs, or other logic circuitry. Control unit 152 may include memory (not shown) that stores computer-readable program instructions that cause control unit 152 to perform the functions ascribed to it herein. The memory may include any magnetic, optical, or electrical media, such as a RAM, ROM, hard disk, CD-ROM, or EEPROM. Control unit 152 may maintain QoS information 120 in the memory.

[0113] FIG. 10 is a flowchart illustrating an example method in which a SE router 112 provides QoS information to a CPE device 114 consistent with the principles of the invention. In particular FIG. 10 illustrates an example method in which the SE router 112 and the CPE device use QoS information to provide packet transmission according to a requested QoS class for a unicast packet flow, which, in the illustrated example, is a VoIP call. When a subscriber using a subscriber device 118 initiates a VoIP call, SE router 112 receives a VoIP request message from the subscriber device 118 (160). The request message may request authentication and authorization for a VoIP call with packet transmission according to a particular QoS class.

[0114] SE router 112 checks authentication/authorization information 142 to authenticate and authorize the subscriber (162), and to retrieve QoS information for the VoIP call from QoS profiles 116 (164). As described above, the QoS profiles 116 may include information describing a route or packet flow that may be used by SE router 112 to provide packet transmission according to the requested QoS class for the VoIP call, and information describing preferential queuing that may be used by CPE device 114 to provide packet transmission according to the requested QoS class for the VoIP call.

[0115] SE router 112 sends a control message to CPE device 114 that includes the QoS information used by CPE device 114 to provide packet transmission according to the requested QoS class for the VoIP call (168). As described above, the control message may be an in-band message, and may be sent to CPE via a dedicated control VC or VLAN. Based on the information contained in the control message, CPE device 114 dynamically configures a QoS profile 120 for a layer-2 link between CPE device 114 and the subscriber device 118 (172). CPE device 114 forwards VoIP packets to the attached subscriber device and to switch via the layer-2 link to provide packet

transmission according to the requested QoS class by, for example, preferentially queuing the VoIP packets (174). SE router 112 forwards VoIP packets to provide packet transmission according to the QoS class indicated by QoS information 116 by, for example, forwarding the VoIP packets on a route or packet flow across network 122 that is designated for VoIP packet traffic (176).

[0116] FIG. 11 is a block diagram illustrating an example multimedia networking environment 180 in which a SE router 182 controls packet forwarding by a switch 184 and a CPE device 186 to provide multimedia service to a subscriber according to an associated service profile consistent with the principles of the invention. A service profile for a subscriber may include, for example, a one or more general QoS classes for packet flows originating from or destined for a subscriber device 188 associated with the subscriber. The service profile may identify, for example, routes or packet flows through a network 190 that SE router 182 may forward packets originating from subscriber device 188 on. The service profile may also identify layer-2 links, e.g., VCs, VLANs, or the like, configured between SE router 182, switch 184 and CPE device 186, that packet flows originating from or destined for subscriber device 188 may be forwarded on. The service profile may identify classes of packets that may be forwarded on preferential packet flows, VCs, VLANs, or the like. Further, the service profile may identify a preference level for queuing of packets originating from or destined for subscriber device 188.

It is clear that what Kokot discloses is that, as defined by a predefined subscriber profile, some particular packet flows may be treated preferentially and may have a QoS class associated therewith (“[t]he service profile may identify classes of packets that may be forwarded on preferential packet flows”). However, again, this does not mean that a client is assigned to a QoS class, as is expressly claimed. At most, any QoS class mentioned in Kokot attaches to a packet flow, and is not assigned to the client itself, regardless of whether the packet flows come from a particular subscriber or not. Thus, Kokot, in fact, does not remedy the deficiencies of Colby.

Also, it is important to point out that Kokot is directed to (see Abstract) a network layer device that controls provision of data link layer functionality by a data link layer device to provide a requested multimedia service to a subscriber. For example, the network layer device may control the performance of multicast elaboration by the data link layer device, or the queuing and forwarding of packets by the data link layer device to facilitate transmission of packets according to a Quality of Service class. The network layer device may send control messages to the data link layer device to

dynamically configure a control object stored by the data link layer device, such as multicast filter information or a Quality of Service profile. The network layer device may be a service edge router, and the data link layer device may be a customer premises equipment device, e.g., a modem or wireless access point, or a switch, e.g., a digital subscriber line access multiplier.

It is unclear why one of ordinary skill in the art, given the HTTP flow switch in the web client/server environment of Colby, would look toward a non-analogous reference that discloses techniques dealing with VoIP call between two subscribers. Thus, Applicants assert that Colby and Kokot are not properly combinable. Assuming for argument sake that the final Office Action is correct in asserting that Colby refers to request submission scheduling (which it clearly does not for at least reasons given above), there is no legally-sufficient motivation that appears in either Colby or Kokot to take any parts of Kokot's subscriber-to-subscriber VoIP call processing steps and integrate them with any scheduling performed in Colby's HTTP request based client/server system.

Since Colby and Kokot are used in each and every other obviousness rejection raised by the final Office Action, Applicants respectfully assert said various obviousness rejections are deficient for at least the same reasons as given above.

In addition, Applicants, after considering the present Office Action in its entirety, respectfully assert the same deficiency arguments presented in their previous response dated May 5, 2008 (the disclosure of which is incorporated by reference herein) with respect to Bender, Veres, Menditto and Lu.

With further reference to Bender, the final Office Action asserts that Bender teaches the step of claim 2 (and also recited in claim 18 in a different form) of withholding the request from submission to the at least one server when the request originated from a client assigned to a first QoS class to allow a request that originated from a client assigned to a second QoS class to meet a response target associated therewith. However, as explained at column 5 of Bender:

At step 108, once the deadline for each uncompleted job is calculated, server system 10 schedules the jobs in accordance with an earliest deadline first ("EDF") methodology. With an EDF methodology, the first job that server system 10 schedules is the job which has the earliest deadline, as found in step 106, relative to all of the other jobs. It then chooses the job with the next earliest deadline, and schedules it second, and so on until all of the jobs have been scheduled.

At decision step 110, server system 10 inquires whether each and every one of the jobs have completion times which is earlier than each job's respective deadline, as found in step 106. If any job is not able to be completed prior to its deadline, then the estimated stretch value is not feasible and is therefore adjusted at step 112. From step 112, the feasibility of the adjusted stretch value is re-checked by returning to step 106.

Thus, Bender schedules a job based on completion times and deadline times associated with *that particular job*. Bender does not withhold the request from submission to the at least one server when the request originated from a client assigned to a first QoS class to allow a request that originated from a client assigned to a second QoS class to meet a response target associated therewith. That is, Bender does not schedule jobs based on a “response target” associated with *a particular QoS class*.

Applicants assert that the various dependent claims are not only patentable for the reasons given above but also because one or more of said claims recite separately patentable subject matter.

In view of the above, Applicants believe that claims 1-25 are in condition for allowance, and again respectfully request withdrawal of the various remaining rejections.

Respectfully submitted,

/William E. Lewis/

Date: February 3, 2009

William E. Lewis
Attorney for Applicant(s)
Reg. No. 39,274
Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, NY 11560
(516) 759-2946